

## **eBay - STAYING ONLINE - ALWAYS**

Case code- **ITSY015**

Published-**2002**

"eBay's persistent site failures stem from a lack of coordinated IT planning and a centralized database and storage structure that creates a single point of failure."

***Internetweek.com, January 11, 2001***

### **THE PROBLEMS**

In January 2001, eBay, the largest online auctioneer in the world, saw a major outage of its website, which lasted 11 hours. Company sources blamed the mishap on some problems with the storage hardware and database software. eBay CEO, Meg Whitman, blamed both the primary and backup infrastructure of the website. To add to the company's problems, it had to delay replacing some of its hardware due to the busy holiday season.

The outage led to significant financial problems for eBay. The company had to extend pending auctions by 24 hours and refund the fees paid for all auctions that were scheduled to end during the outage period. According to an analyst <sup>[1]</sup>, the outage cost eBay \$450,000 in lost service fees - about 0.3% of its total revenue for the first quarter of 2001. The next day, eBay's stock went down by 8.85% to \$35 7/8. This was not the first instance of eBay's outage. The company had seen 15 similar outages between August 1998 and November 1999, including a major one in June 1999 that made the site inaccessible for 22 hours. During the two-day outage, eBay's stock came down from \$182 to \$135 reducing its market capitalization by \$5.7 billion. The outage was prolonged as database files became corrupted and files had to be rebuilt before the system could be brought online.

Experts said that eBay had failed to build a redundant storage hardware and scalable web architecture. They added that eBay had been functioning largely as a marketing-driven company, which placed far more importance on adding new features <sup>[2]</sup> rather than having a core infrastructure to avoid outages.

### **BACKGROUND NOTE**

eBay's founder Pierre Omidyar (Omidyar) felt that an Internet auction site could create a highly efficient market. The idea seemed to have originated out of the difficulty that his fiancée faced while trying to collect Pez dispensers from the San Francisco Bay area. Omidyar launched Auction Web (the site's domain name was ebay.com) in September 1995 to enable people to trade.

eBay started as a free website and quickly attracted so much traffic that Omidyar needed additional infrastructure to keep pace with the growing business. The company started charging a nominal fee of 10 cents for listing an item and also collected a percentage of the final sale value from the sellers. In September 1996, Omidyar brought in Jeff Skoll, a Stanford graduate, to manage the business, while he concentrated on the technological aspects.

As the interaction among individual buyers and sellers increased, complaints from buyers and sellers began to mount. Omidyar urged them to work things out amicably among themselves. He said, "If traders are going to complain about people they do not like, they should be willing to say something nice about people they do." In February 1996, Omidyar started a virtual 'feedback forum' where customers could share their trading experiences

In February 1998, Meg Whitman <sup>[1]</sup>, a Harvard MBA, joined eBay as its President and CEO. Whitman initiated efforts to arrange an Initial Public Offering (IPO) for eBay. In September 1998, eBay

offered 4.025 million shares to the public at \$18 per share. It received a tremendous response and the stock closed at \$47 on the first day of its listing. EBay's main product categories were automobiles, antiques & art, books, movies, music, coins & stamps, collectibles, computers, dolls, doll houses, jewelry, photo & electronics, pottery & glass, sports, toys, premium arts & collectibles and miscellaneous items.

Its operations included registration of users, facilitating the auction and maintaining relevant databases. The company provided seven days availability with a few hours maintenance period each week. Members were notified by e-mail when they registered on the site. The company notified them when they were either outbid or had won an auction. The company also sent daily status reports to active sellers and bidders regarding their current auctions. It also maintained a database regarding member registration information, billing accounts, current auctions and historical listings. Customers could use the relevant links to get information regarding how to use the site. They were also informed about the tools that eBay had to ensure the site's safety.

A bidder could see the listing of the total items to be auctioned by clicking categories or specific sub-categories. A 'shop by photos in the gallery' facility allowed a bidder to see the photograph of the items along with the description. Users who wanted to trade in the local region could specify region in the text box (pick a region) and could begin trading by choosing a suitable category (automobile, appliances). A seller could get information regarding activities like registration, how to enter the credit card number and how to photograph an item. Clicking the 'sell' button, a seller could also visit the page where he could list the item for sale.

In addition to a search facility on its home page, eBay provided a 'search' link on the main navigation bar to enable visitors to search items in six ways; listing title by search, item number, seller search, bidder search, completed auction and regional search. Search facilities also provided the option to narrow down the number of items to be listed after the search. For instance, while searching items through 'listing title search,' one could specify the price range and country where the item was available and could sort the listing based on the last date of auctions.

eBay also sent electronic invoices to all sellers via e-mail on a monthly basis. The accounts of the sellers, who opted for credit card payment mechanism, were billed directly. At regular intervals, eBay updated a text-based search engine with the titles and descriptions of new items, as well as the pricing and bidding updates for active items. The company also supported a number of community bulletin boards and chat areas where its customer support personnel could interact.

By 2001, eBay had emerged as the largest person-to-person auctioneer in the world with a registered user base of 42.4 million. For the financial year ended December 31, 2001, the company generated \$748.82 million revenue and \$90.45 million net profit. eBay was ranked as the #1<sup>[2]</sup> online shopping destination, with nearly 5 million unique daily visitors during the peak shopping season in December. For the quarter ended December 31, 2001, eBay hosted over 126.5 million items. The company's users transacted a record \$2.735 billion in gross merchandise sales during the same period.

Though the company had been performing well financially, the frequent outages had attracted severe criticism. According to a former staff member, eBay had not installed redundant storage hardware<sup>[3]</sup> to take over if the main storage system failed, even after six years of the site's launch in 1995.

Though the management had been warned about the risk of storage failure and was recommended an upgrade in the storage infrastructure, the expense was not approved. eBay's initial applications relied on a single large database. Till January 2001, a single Oracle database supported the item listings and auction operations. When eBay began to split the database, resource constraints and a divided management came in the way of initiating major changes

## **TECHNOLOGY**

EBay's infrastructure, built in 1995, was initially designed in-house to support buying and selling between small groups of collectors and hobbyists. The company used Pentium based Internet servers that ran on Windows NT operating system. For Internet connectivity, it had entered into a partnership with Exodus and Above Net Communications located at Santa Clara in California. These

two companies hosted eBay's web servers, database servers and Internet routers. Most of eBay's software had also been developed in-house.

After the June 1999 outage, eBay realized that its infrastructure was not flexible, scalable and reliable. It initiated steps to strengthen the server infrastructure. The company decided to replace the existing Sun enterprise 10000 servers, working singularly and with limited scalability, with IBM mainframe and AIX Unix servers. In the first phase, the company implemented a backup solution that allowed the network to recover within four hours by increasing redundancy on servers, routers<sup>[1]</sup>, switches and RAID<sup>[2]</sup> drives. The company then built another back up solution - a running duplicate of major systems that reduced the recovery period to one hour. In the second phase, eBay set up an eight-person IT group to evaluate and implement a more sophisticated network..

A group's main challenge was to ensure the database's resiliency, which had more than three million items appearing concurrently. As the entire list was in a single database, even a minor corruption could affect the entire network. The existing hardware was also fast approaching a saturation point. To overcome these weaknesses, eBay decided to create separate databases for different auction categories. This ensured that in case of corruption, only one database would go down and only the people participating in that related category would be affected by the outage.

In November 1999, eBay tied up with Sun Microsystems, VERITAS and Oracle to strengthen its hardware infrastructure. The solution was based on tightly integrated clusters of Sun's enterprise 10000 servers, Storage disk arrays<sup>[3]</sup>, Solaris Operating Environment<sup>[4]</sup>, Oracle 8i enterprise relational database<sup>[5]</sup> and a suite of VERITAS products.

Before the solution was implemented eBay's server environment was thoroughly evaluated, including hardware inventories, configurations, performance, resource consumption and future growth potential. To capture the peaks and troughs of system performance due to the highly fluctuating traffic, statistics related to CPU utilization, CPU wait input/output, disk service time, disk utilization and memory usage were collected over several days to reflect varying workloads. The data provided a detailed system performance report that resulted in a decision to go for the new cluster configuration.

The hardware solution consisted of a two-tiered, clustered architecture. The top tier consisted of one enterprise 10000 server, (also known as Star fire) that drove e-commerce and VERITAS front-end database applications such as Cluster Server, Database Edition for Oracle, Netback up, Volume Manager and File System. In the bottom tier, another enterprise 10000 server powered Solaris Operating Environment and an Oracle8i database. Maynard Webb, President of eBay Technologies, explained the benefits of this architecture, "The two Star fire servers supporting eBay's production environment are clustered and a third provides failover for disaster recovery. The cluster gives us the ability to quickly failover: if a server or database is unavailable, the other can take over"

Another Star fire server was used to support eBay's quality assurance operations and two for the company's test environment. Enterprise 6500 and Enterprise 4500 servers powered the search engines. Additionally, Storage disk arrays housed over 2.7 terabytes of data. Numerous workstations also supported various functions including software development and system administration. The enterprise servers provided an ideal platform to eBay for numerous applications including enterprise resource planning, electronic commerce, data warehousing, Internet/Intranet and customer management systems. With the clustered server architecture, eBay significantly improved its site availability. The company had also reduced the average downtime of failover to another machine, which was previously between two to four hours to between 10 minutes and one hour.

The new hardware solution helped eBay to handle much higher traffic compared to its previous setup. By late 2000, eBay's site was handling about 130 million page views with more than 18 million registered users compared to 65 million page views and about 8 million registered users in June 1999. The site supported about 18 million searches a day compared to three to four million searches a day previously. The company also managed to lower its operating cost by two percent in 2000. However, despite the above infrastructure, eBay faced another outage in January 2001. A

hardware failure in one of the backup systems had caused the outage. After the system was restarted, another problem in the primary and backup systems brought down the site.

eBay's third backup system was not of much help as a database problem brought down the site again after 40 minutes. Industry experts felt that though the clustered server approach was easier to administer, it was also more prone to breakdowns as the servers shared central resources, like the operating and storage systems. They argued that other high-traffic sites used more loosely integrated server farms<sup>[1]</sup>, in which many small, low-end servers perform identical functions in a redundant configuration. The search engine Google, for instance, used 4,000 Linux servers that searched and served web pages in parallel, while Yahoo used clusters of Unix servers. If one server went down, the rest picked up the load.

To tackle the problems in its storage system, eBay partnered with Sun to develop a storage solution, with high availability, resilience, and serviceability known as Storage Area Network (SAN)<sup>[2]</sup>. Webb explained the benefits, "We have been moving and segmenting our entire database server into smaller servers. This helps ensure we don't have any single point of failure. In other words, an outage will affect just a portion of the system not the whole site. It will also be more cost-effective because the SAN will enable us to have six or more front-end machines to every two-recovery machines, instead of a one-to-one cluster. A SAN will also allow us to scale the system without accruing any downtime."

In September 2001, eBay went on a major software revamp to improve its site availability and make it more dynamic and interconnected. After reviewing more than 20 vendors, the company finally opted for IBM and its Web Sphere application server and began to put in place its 'V3' application architecture. eBay also revamped its server-side application development architecture to support the Java 2 Enterprise Edition (J2EE)<sup>[3]</sup> and Enterprise JavaBeans<sup>[4]</sup>. The company replaced a C++<sup>[5]</sup> object framework that required a lot of structural programming. Though the C++ environment was more flexible than COBOL, it could not compete with J2EE, which was becoming the de facto application development framework. J2EE's objects could handle a much higher level of abstraction than C++. The new applications were more widely distributed, running across multiple machines (both Windows and Solaris) and relied on data-dependent routing, which utilized server cache<sup>[6]</sup> more efficiently.

The V3 deployment was phased in slowly. It began during the fourth quarter of 2001 and was to be completed over the next 18 months. By December 2001, eBay's IT infrastructure had not only become flexible enough to handle instant recoding but also sturdy enough to process more than 800,000 transactions every minute.

## **REAPING THE BENEFITS**

The fact that auctions were relatively static events (until the final few minutes when lot of bids could often put a major hit on its servers and apps) was the biggest challenge for eBay. According to industry experts, eBay's decision to opt for a more distributed application infrastructure had reduced the load on its infrastructure. Besides, the software developers had more flexibility to change the codes without affecting the entire application. According to an analyst<sup>[1]</sup>, eBay's 99 percent uptime was in line with other consumer sites like Amazon.com and Yahoo. However, it was more important for eBay to ensure zero outages than other consumer sites.

Todd said, "If Yahoo! goes down, I can't get news or e-mail for 20 minutes or an hour, but if eBay goes down, I might miss out on a \$1,000 auction that I am running my business on." The company's efforts seemed to have paid off, as it did not face any major outages after the January 2001 episode. The company posted a handsome growth in earnings and net profits in 2001 despite a global slowdown in the information technology industry. Webb said, "Analysts in our industry were asking, 'Can eBay scale to meet increased demands? Can eBay support its growth?' We do not get those questions any more."

**QUESTIONS FOR DISCUSSION:**

1. Briefly analyze the online auction process as followed by eBay. Explain why it is important for a website like eBay to ensure zero outages.
2. Discuss the steps taken by eBay to strengthen its hardware, software and storage infrastructure to ensure its continual availability. What benefits did the company reap after taking the above measures?